Fermilab

Managed by Fermi Research Alliance, LLC for the U.S. Department of Energy Office of Science

Advancing Particle Accelerator Science with High Performance Computing

James Amundson, Fermilab

CMSE Workshop

2015-09-16

My ultimate topic: Computational Beam Dynamics

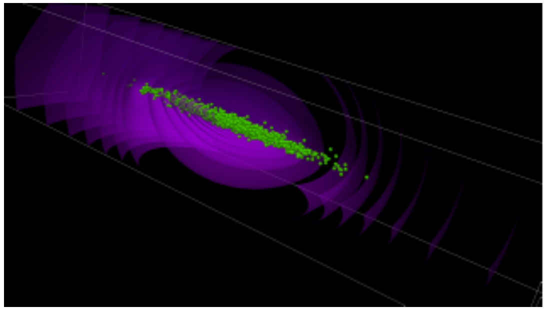

Synergia: A comprehensive accelerator beam dynamics package

http://web.fnal.gov/sites/synergia/SitePages/Synergia%20Home.aspx

Fermilab

Accelerator Simulation Group

James Amundson, Qiming Lu, Alexandru Macridin, Leo Michelotti, Chong Shik Park, (Panagiotis Spentzouris), Eric Stern and Timofey Zolkin

Computer time from INCITE

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

The ComPASS Project

High Performance Computing for Accelerator Design and Optimization

https://sharepoint.fnal.gov/sites/compass/SitePages/Home.aspx

Funded by DOE SciDAC

CAMPA Consortium for Advanced Modeling of Particle Accelerators

Funded by DOE

This is a cross-disciplinary talk

Computational Accelerator Physics

advances

Accelerator Physics

advances

Particle Physics

Particle Physics

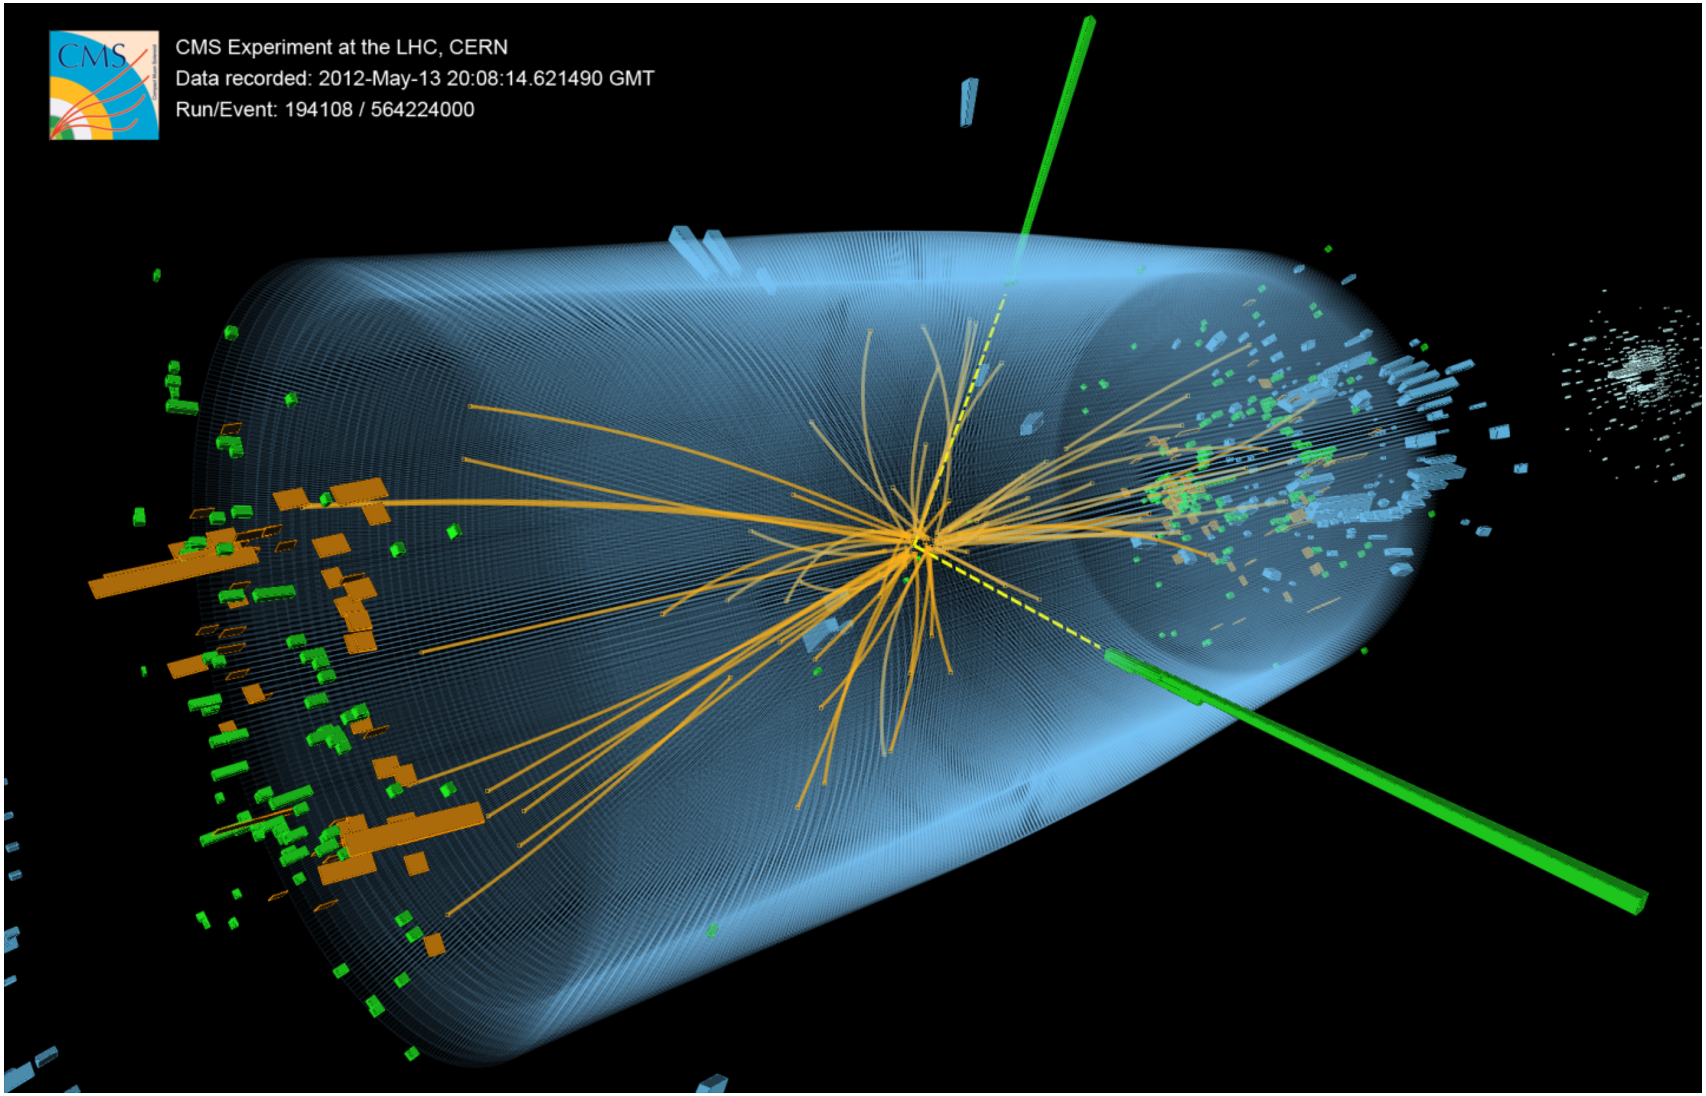

(Accelerator-based) Particle Physics

- Experimental Particle Physics is a statistical science
  - Large numbers of interactions must be analyzed
  - This is the distinguishing feature of the problem domain
- Particle physics experiments include
  - Searching for new particles
  - Measuring the properties of known particles
- Particle Physics is an international effort
  - Fermilab is *the* U.S. national lab devoted to particle physics
  - A significant fraction of our current work at Fermilab is devoted to various aspects of the CERN program, especially in computing

Particle Physics Experiments Require Many Interactions

- The top quark was discovered at Fermilab in 1995.
  - Roughly 1 in 1,000,000,000,000 Tevatron collisions produced a top quark.
- The Higgs Boson was discovered at CERN in 2012,
  - Supporting evidence came from Fermilab (Tevatron).
  - Roughly 1 in 100,000,000,000,000 LHC collisions yielded a distinguishable Higgs Boson.
- Many properties of neutrinos are still to be measured
  - An MeV-neutrino can travel through a light-year of lead with only a 50/50 chance of interacting.

More interactions means greater precision and greater chance for discovery.

Fermilab Accelerator Complex

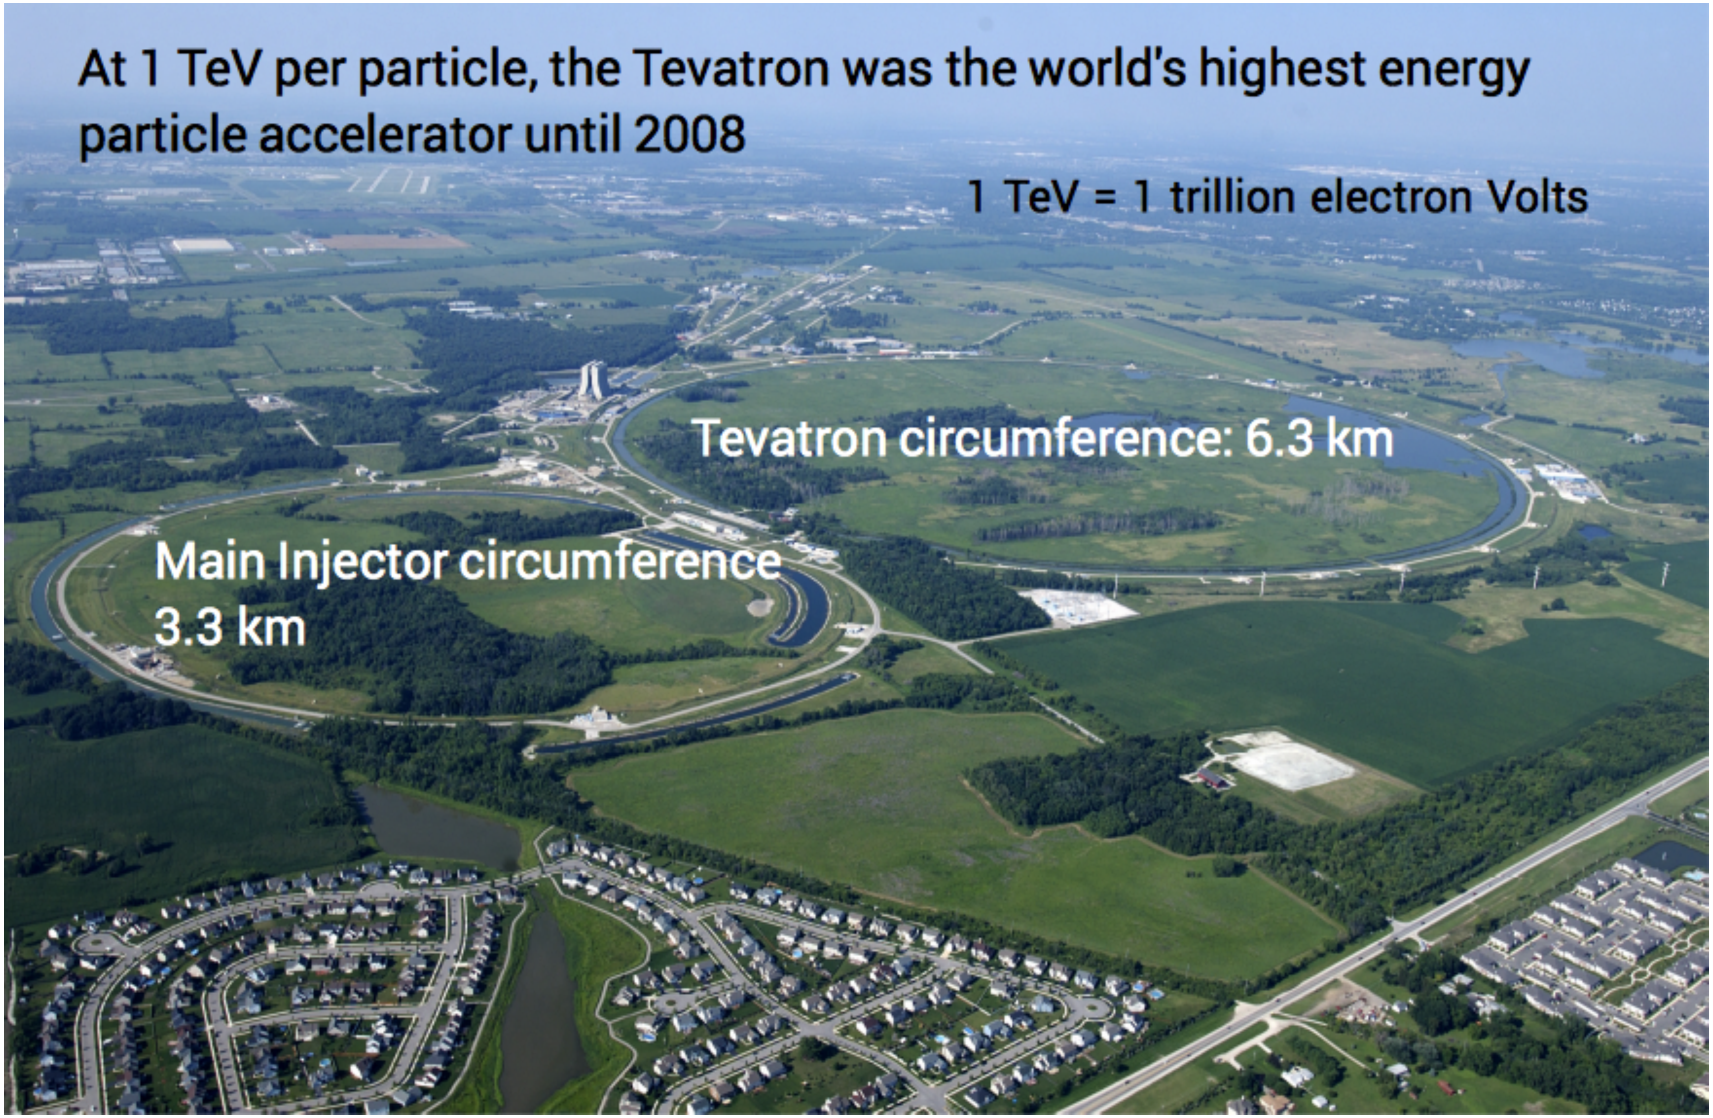

The CERN Accelerator Complex

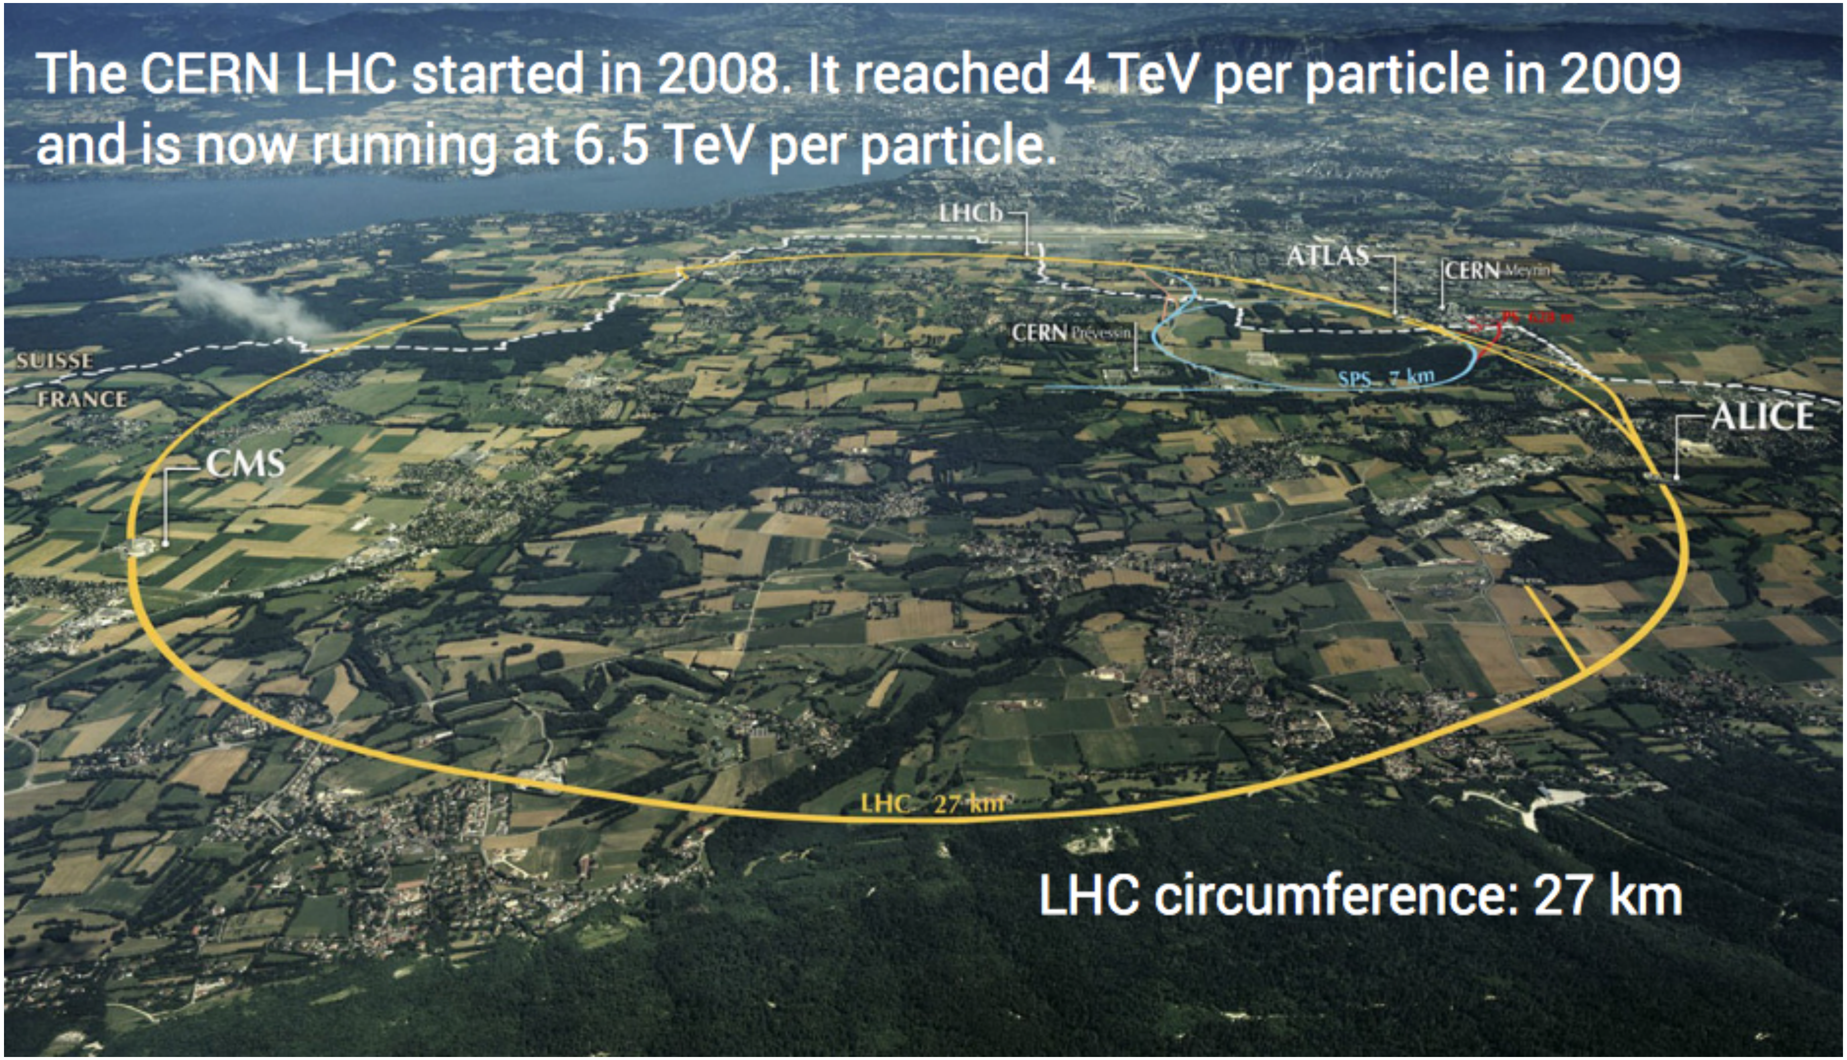

Particle Physics in Transition

- The Fermilab Tevatron collided protons with antiprotons at the highest energy to date
  - The (old) *Energy Frontier*
- The CERN LHC collides protons with protons at the highest energy ever
  - The (new) *Energy Frontier*
- Fermilab is now focusing on providing the world's most *intense* beams, especially neutrino beams
  - The (new) *Intensity Frontier*
- Fermilab plans to create even higher-intensity beams
  - The future of the Intensity Frontier
- CERN plans to increase the intensity of the LHC
  - The future of the Energy Frontier

The Intensity Frontier Now

The Intensity Frontier: Near Future

Searching for new physics at the Intensity Frontier

Fermilab Experiments

- First, g-2: precision measurement of virtual particles in the quantum vacuum

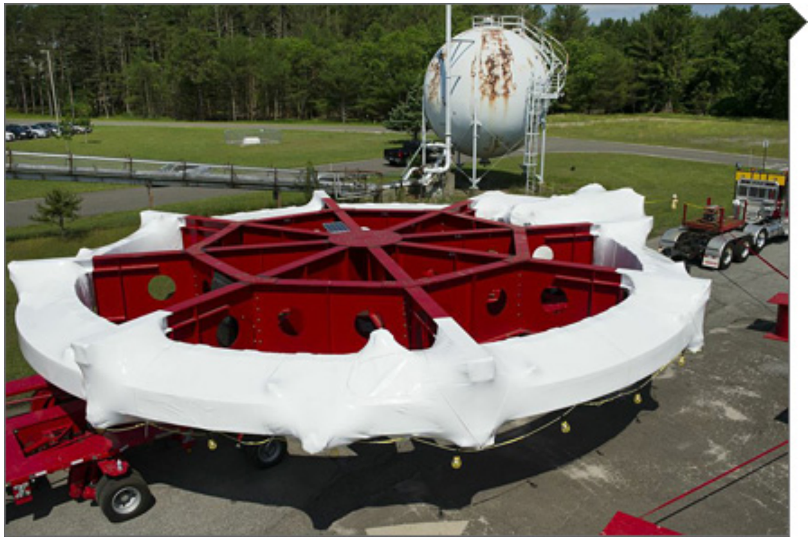

- Next, Mu2e: search for (thought to be forbidden) decays of muons into electrons

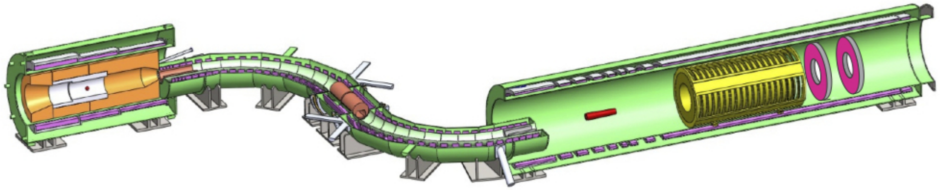

Intensity Frontier: Longer-term Future

- Fermilab will produce even higher-intensity beams (PIP-II)
  - Long Baseline Neutrino Facility (LBNF) and the Deep Underground Neutrino Experiment (DUNE) in South Dakota

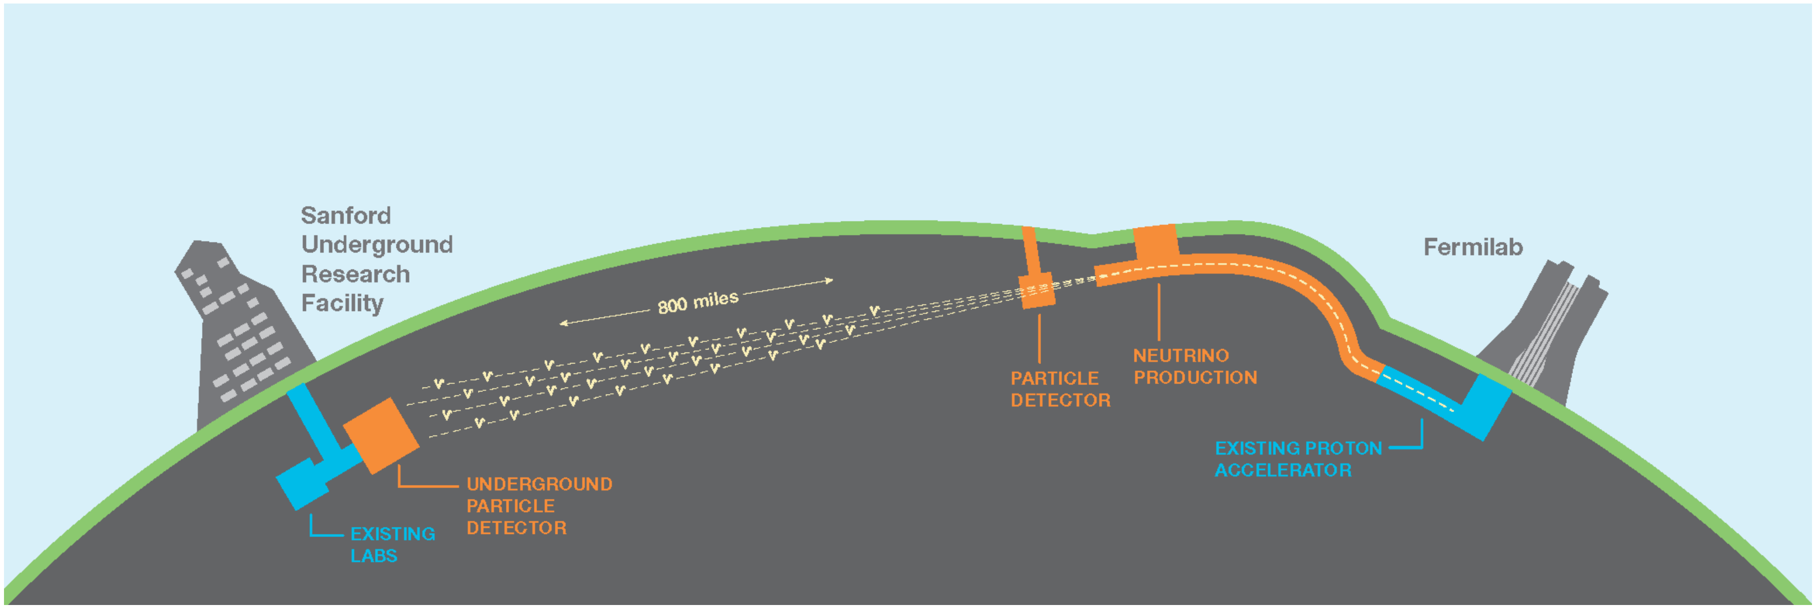

Accelerator Physics

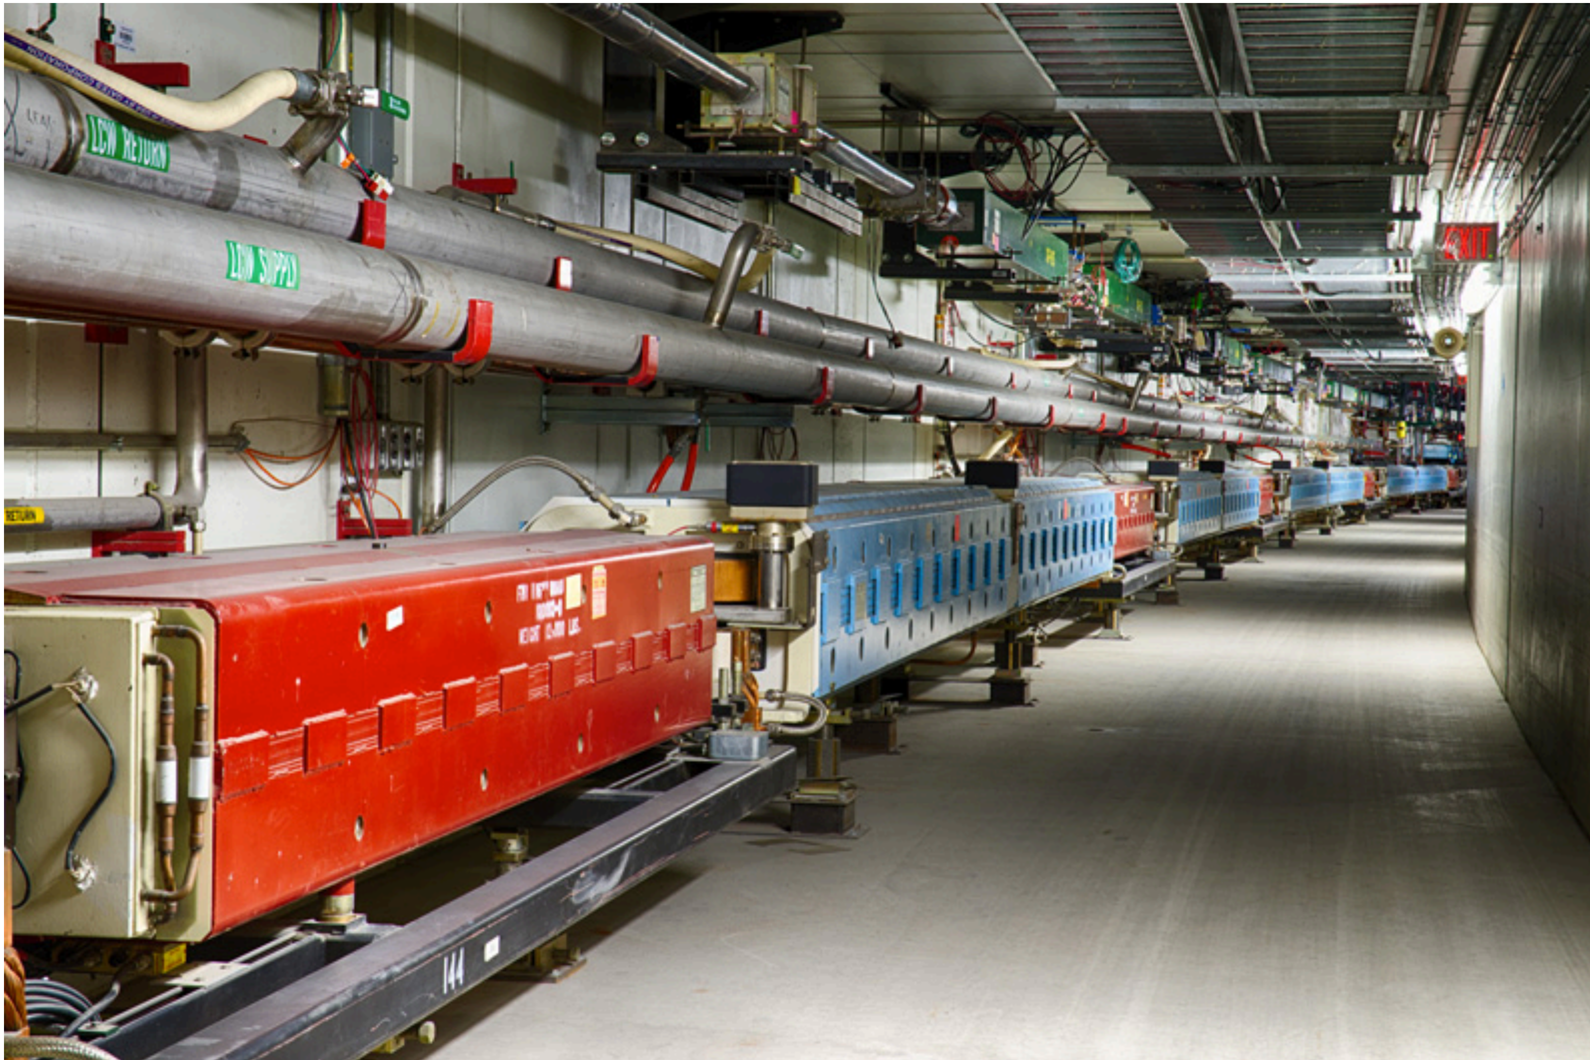

Particle Accelerators

- Particle accelerators enable discovery in basic research and applied sciences
  - Probing fundamental laws of nature, discovering new particles
  - Studying properties of nuclear matter
  - Studying structure of crystals, amorphous materials, and organic matter
  - Enhancing quality of life: medical treatment, nuclear waste transmutation, industrial applications

Numerical modeling and simulation are essential for the development of new acceleration concepts and technologies and for machine design, optimization and successful operation

Topics in Accelerator Physics Modeling

- Electromagnet structure modeling
  - High fidelity simulations of complex structures

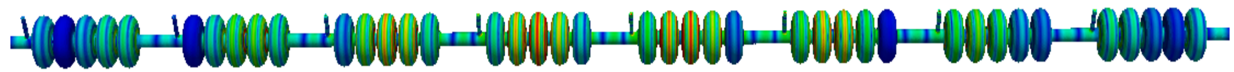

A trapped monopole mode (2.413 GHz) in Fermilab PIP2 650 MHz cryomodule consisting of 8 SRF cavities

- Advanced Accelerators
  - Much higher acceleration gradients than today's technologies
    - Plasma-based
    - Dielectric laser

2-color injection, WARP

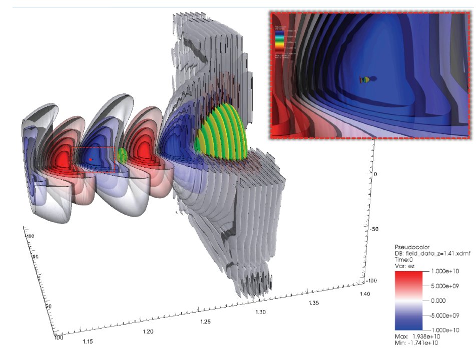

- Beam Dynamics
  - Dynamics in the presence of external and internal fields

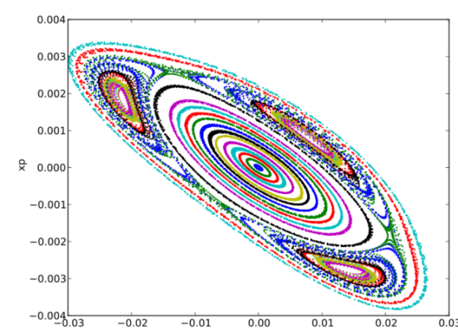

Space-charge trapping experiment in the GSI SIS18 lattice

Beam Dynamics

- Consider the effects on the beam of *all* fields
  - External fields (single-particle effects)
    - Focusing magnets
    - RF accelerating cavities
    - etc.
    - Independent of intensity
    - Dominant forces
  - Internal fields (collective effects)
    - Space Charge: mutual repulsion within a bunch
    - Wakefields: within a bunch and bunch-to-bunch
    - etc.
    - Depend on intensity
      - Set upper limits on accelerator intensity

The push for higher intensity accelerators requires accurate modeling of beam dynamics with external and internal fields.

Computational Accelerator Physics

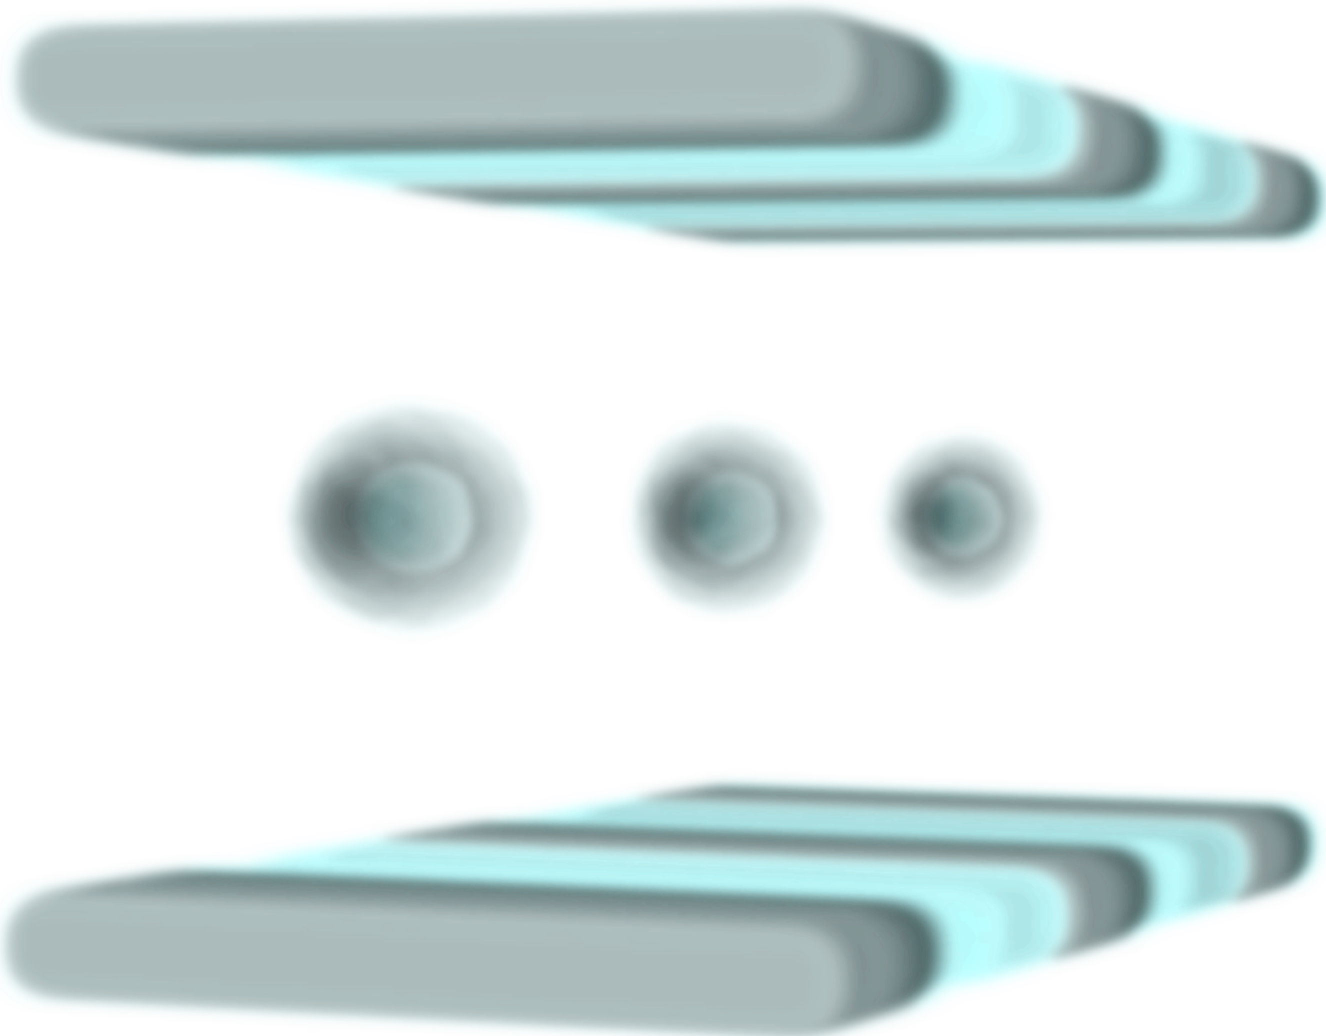

Computational Beam Dynamics

- Existing and planned accelerators
  - 1,000s of elements
    - 10s of *types* of elements
  - 1,000s to 1,000,000s of revolutions
  - 1-1000s of bunches of $O(10^{12})$ particles

- 50-1000 steps/revolution
- Internal and external fields
  - External field calculations trivially parallelizable
  - Internal field calculations require Particle-In-Cell (PIC)
    - Minimal bunch/field structure

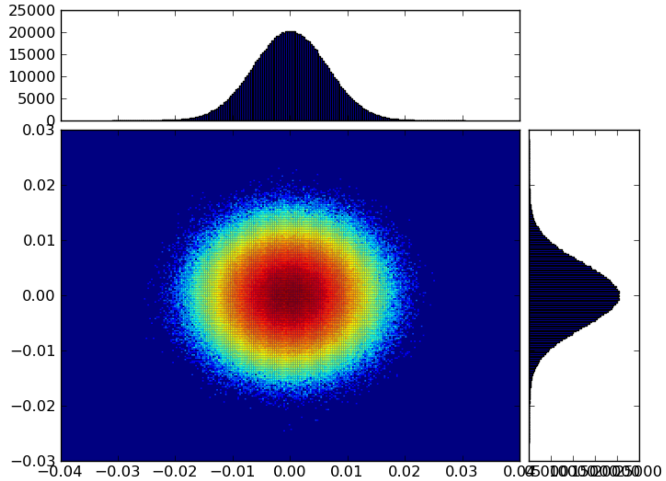

Particle-in-Cell (PIC) Beam Dynamics

- Split operator technique to combine external fields (magnets) with internal fields (space charge, wakefields)
  - 1/2-step external fields + full step internal kick + 1/2 step external fields
- External field calculations (single-particle effects)
  - Trivially parallelizable
- Internal field calculations (collective effects)
  - Deposit particle charge on (local) grid
  - Steps involving communication
    - Combine to form global charge density
    - Solve field equation on grid (in parallel)
    - Send full grid to all processes
  - Interpolate fields to particle position and apply kick

Synergia

- Synergia is a general framework for beam dynamics simulation
  - Emphasis on collective effects
  - Unique ability to perform with true multi-bunch simulations
- A mix of C++ and Python
  - all computationally-intensive code is written in C++
  - user-created simulations are usually written in Python
  - pure-C++ simulations are possible
- Synergia provides a set of functions and classes for creating simulations
  - many examples available
- Virtually every aspect of Synergia is designed to be extendable by the end-user
  - code in C++ and/or Python

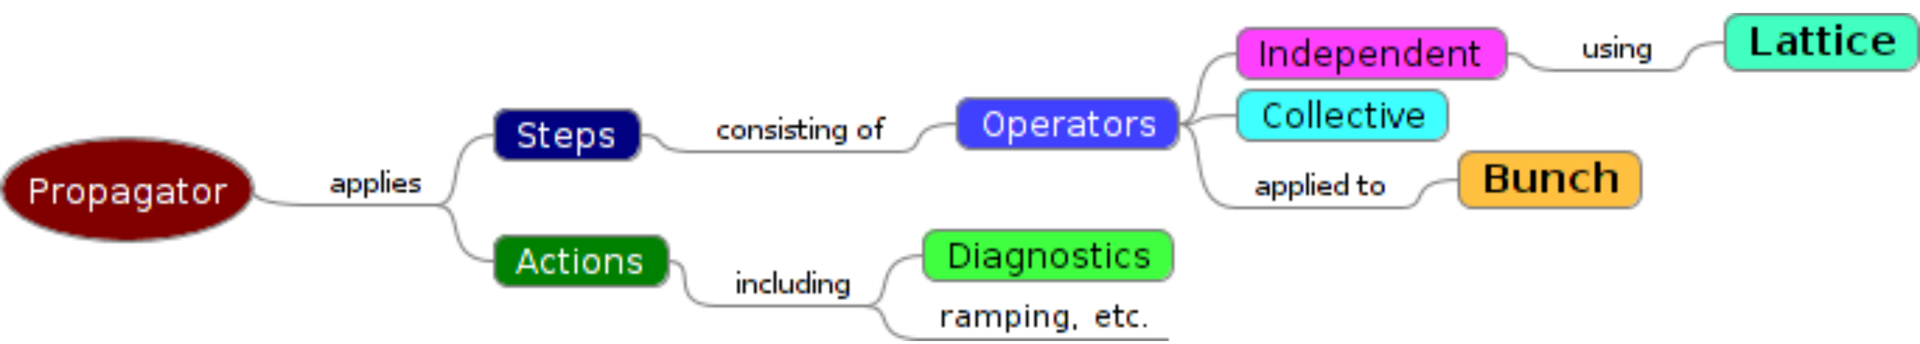

Fermilab

Synergia Runs on a Large Variety of Platforms

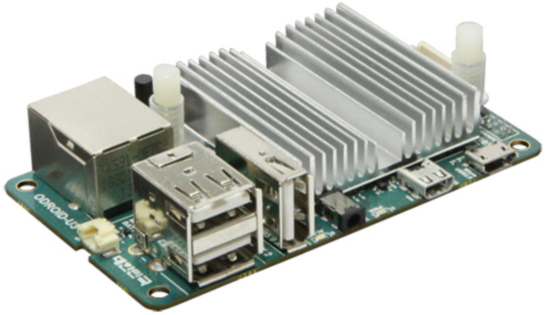

Odroid U3 (Arm A9)

Linux Clusters

Cray

Laptops and Desktops

Blue Gene

A Synergia simulation

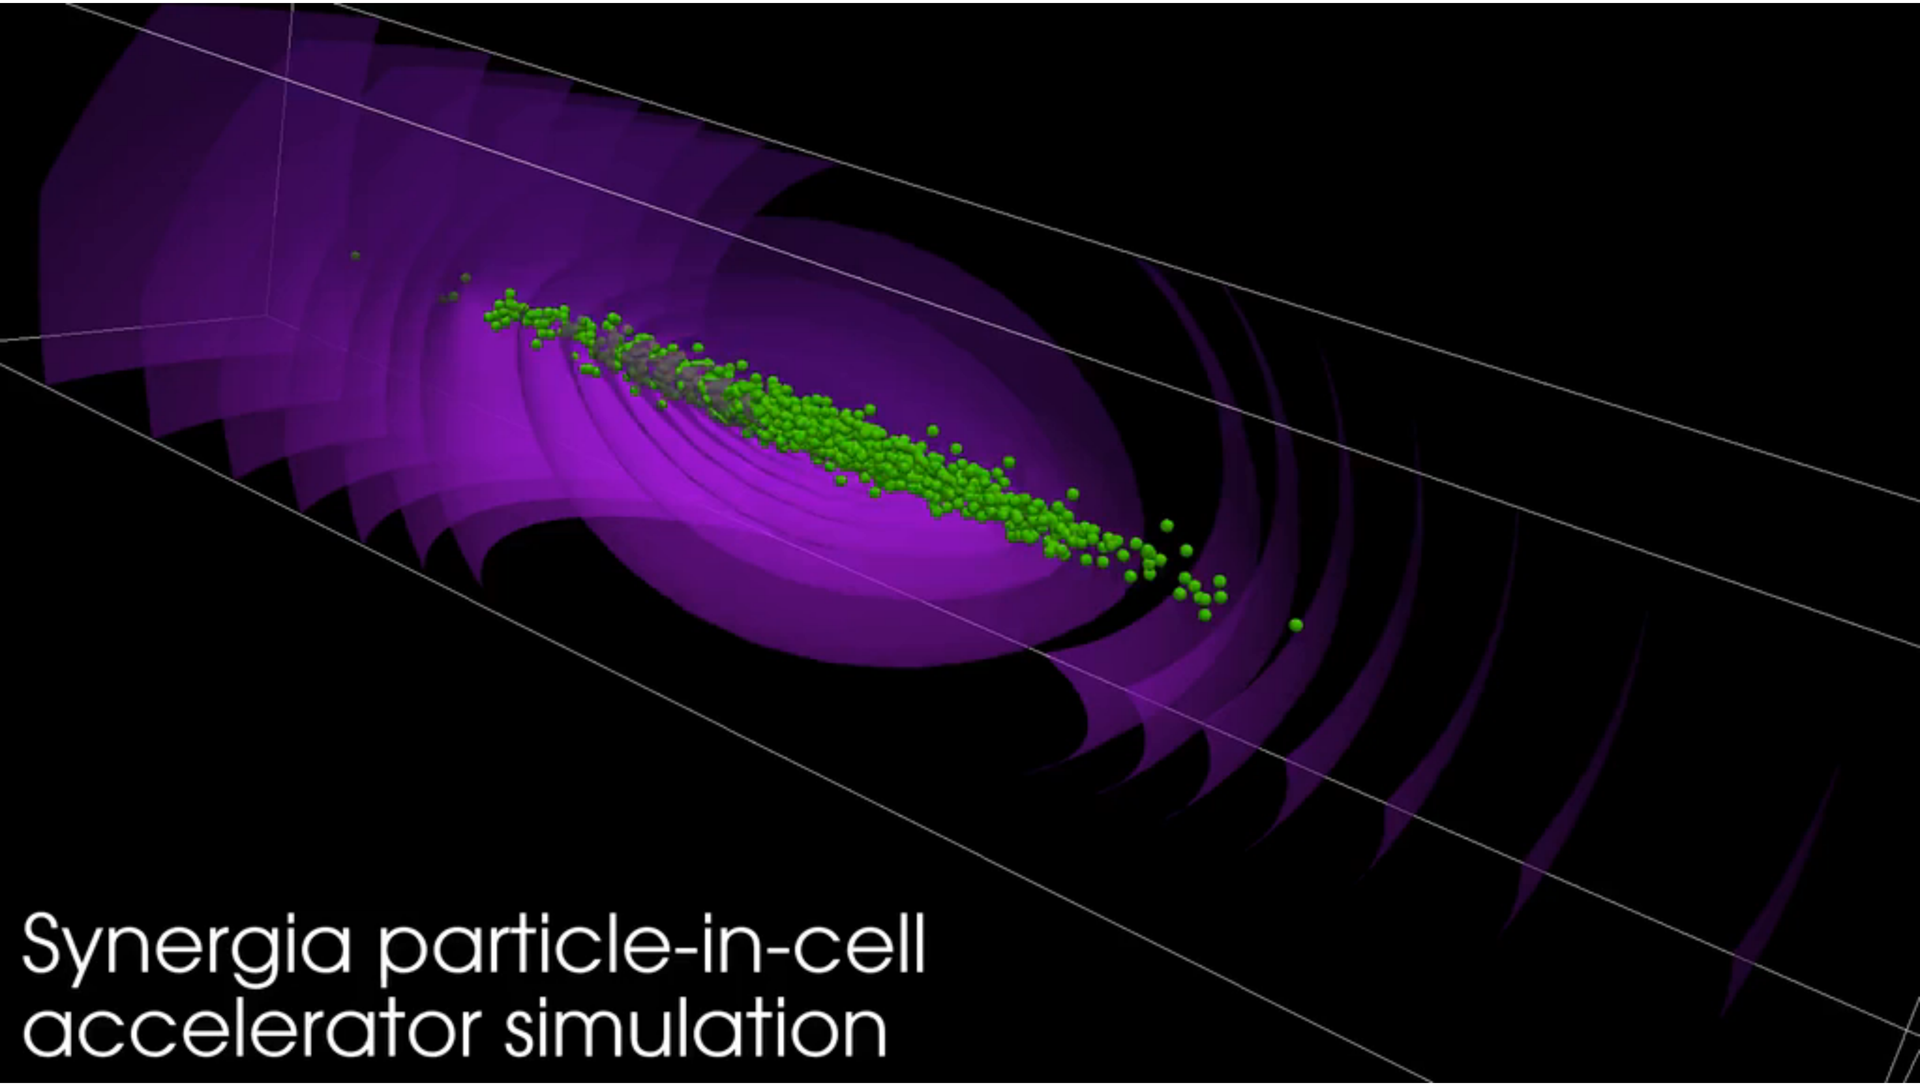

Parallel Scaling in Synergia

- Challenge: beam dynamics space charge simulations are big problems require many small solves
  - Typically 64^3 – 128^3 grids (2×10^5 – 2×10^6 degrees of freedom)
  - Need to do many time steps (10^5 to 10^8)
- Typical pure-PIC scaling applies to scaling with respect to grid size
  - Including decomposing particles by grid location
  - In beam dynamics, external fields can cause particles to move over many grid cells in a single step
    - Communication required to maintain decomposition and load balance
      - Point-to-point communication
      - Complicated for both programmer and end-user
      - Change in physical parameters can change comm. time by x100

Particle (lack of) Decomposition in Synergia

- First step: eliminate particle decomposition
  - Requires collective communication
    - But not point-to-point
    - Collectives are typically highly optimized
  - Simpler for programmer and end-user

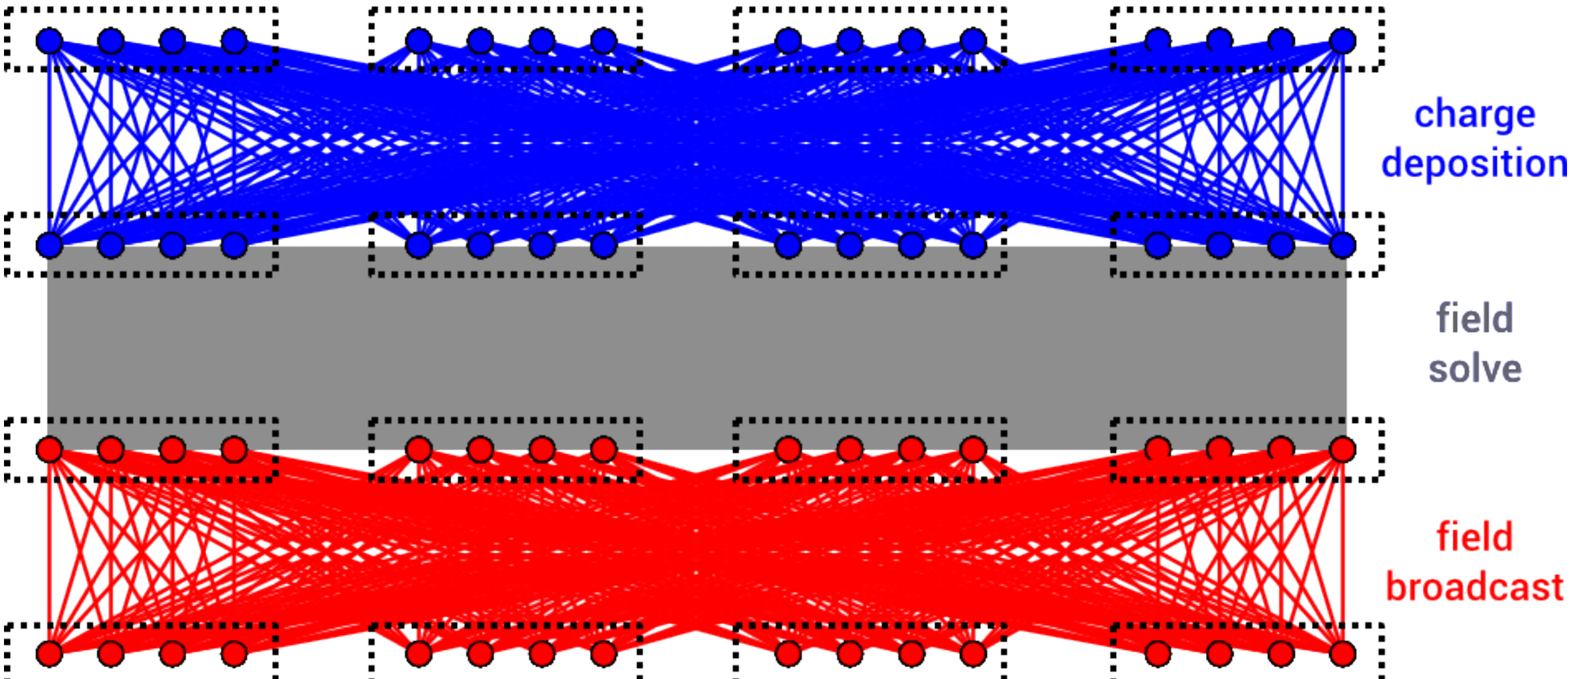

Communication Avoidance

- Second (breakthrough!) step: redundant field solves
  - Field solves are a fixed size problem
  - More calculation, less communication
  - Allows scaling in number of particles and/or bunches
  - Can use arbitrary unit size, but one node is usually best

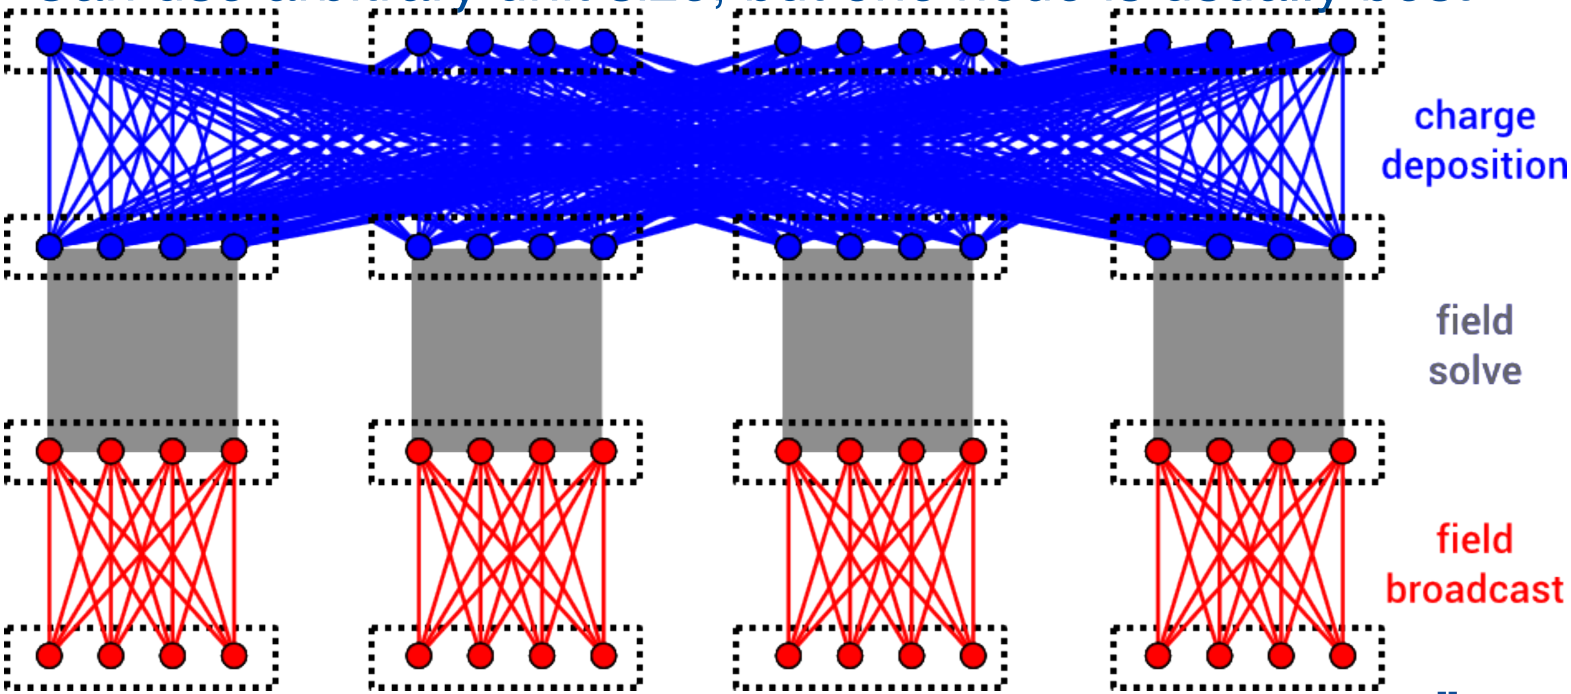

Synergia Scaling with Communication Avoidance

- Synergia
  - Single- and multiple-bunch simulations

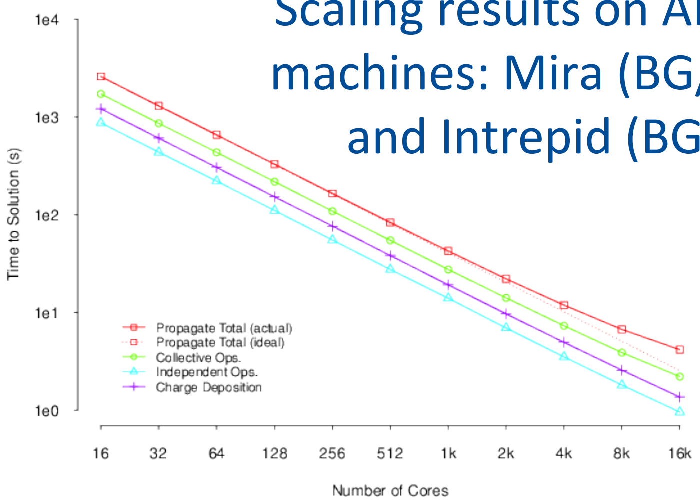

Scaling results on ALCF machines: Mira (BG/Q) and Intrepid (BG/P)

Single-bunch strong scaling from 16 to 16,384 cores

32x32x1024 grid, 105M particles

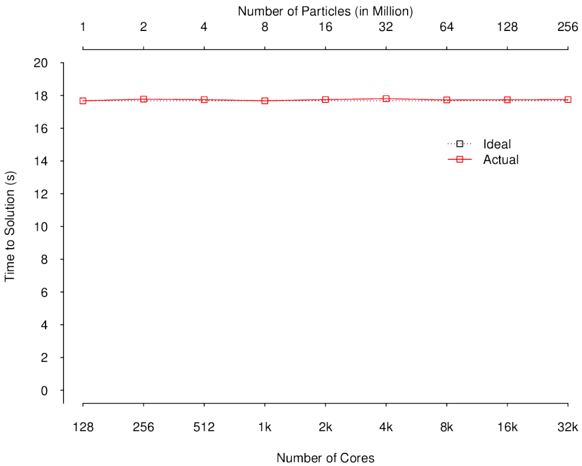

Weak scaling from 1M to 256M *particles* 128 to 32,768 cores

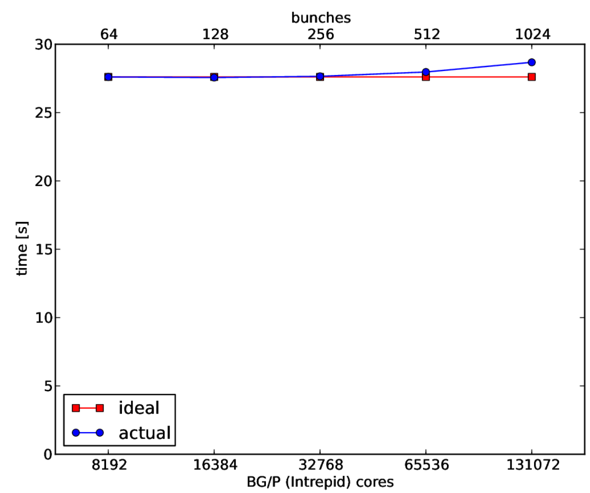

Weak scaling from 64 to 1024 *bunches* 8192 to 131,072 cores Up to over 10^{10} particles

Hierarchical Parallelism in Synergia Simulations

1. Many particles/many grid points
2. Redundant solves within a bunch
   - Communication avoidance
3. Many bunches
4. Parameter scans/optimization

Four levels of parallelism give us the potential to scale to a huge number of parallel processes.

A few recent Synergia applications

Applications

Modeling the Fermilab Booster

- Rapid cycling synchrotron
  - Over 40 years old
- Current intensity ≈ 4.5 x 10^{12} protons per batch
- 400 MeV injection energy
- Observe instability and beam loss at high intensity
  - Space charge important
  - Wake fields also important
- Future Fermilab program will require higher intensities

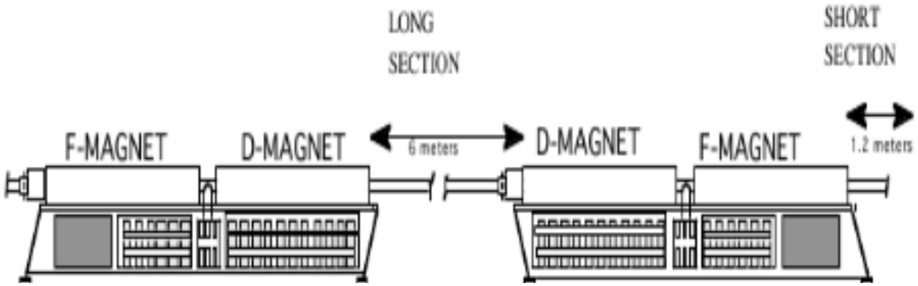

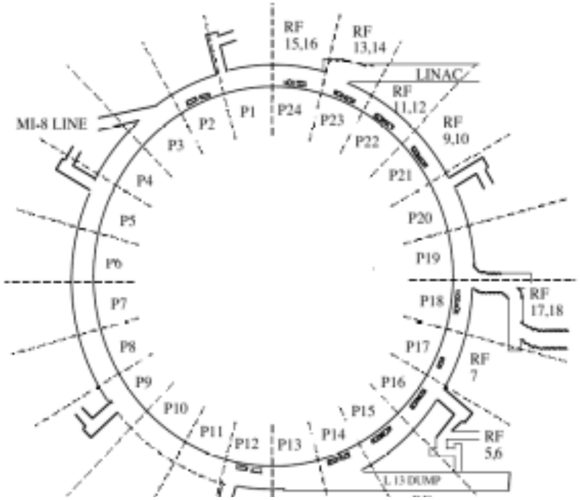

Simulating an Observed Instability

experiment

4 x 10^10 p per bunch $\left(\frac{\omega_{\xi x}}{\beta c}, \frac{\omega_{\xi y}}{\beta c}\right) = 2\pi \times \left(0.06\, m^{-1}, 0.025\, m^{-1}\right)$

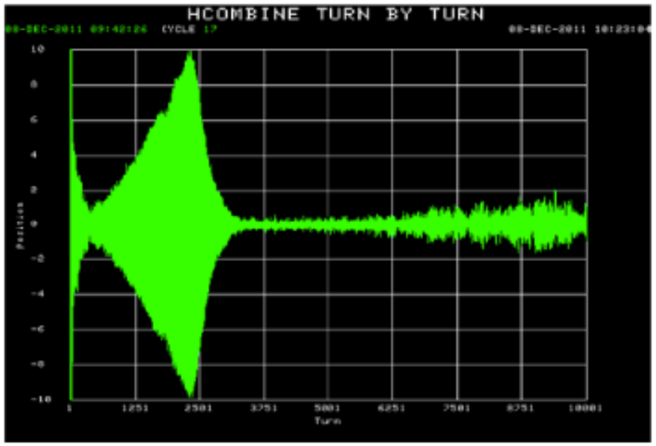

Figure 1: Combined TBT signal from HBPMs (arbitrary units) at $N_p = 4 \cdot 10^{12}$ after coupling correction.

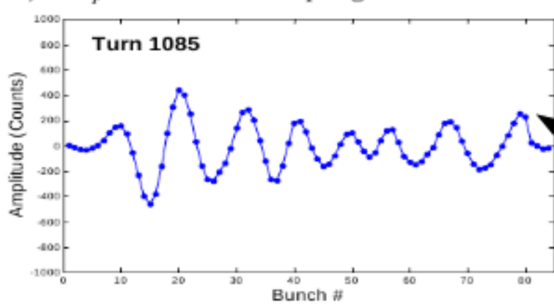

Figure 3: Bunch-by-bunch horizontal positions at the onset of horizontal instability

simulation

5 x 10^10 p per bunch $\left(\frac{\omega_{\xi x}}{\beta c}, \frac{\omega_{\xi y}}{\beta c}\right) = 2\pi \times \left(0.023\, m^{-1}, 0.023\, m^{-1}\right)$

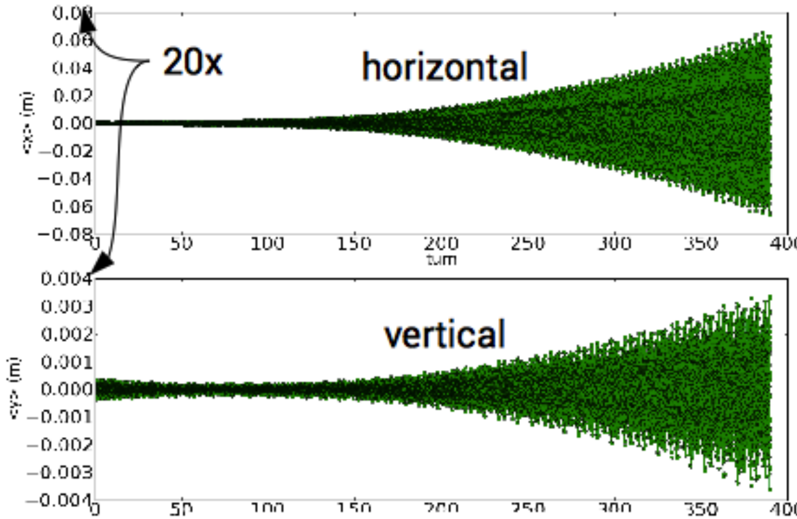

bunch-to-bunch coupling effect

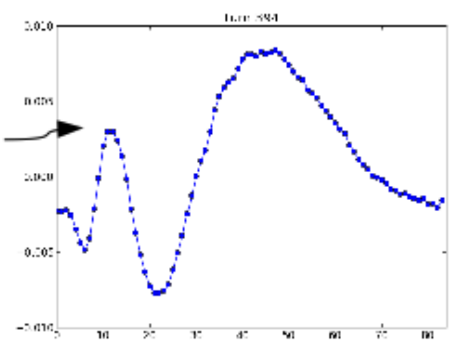

qualitative agreement

Fermilab

Why? Multi-bunch Instability

6 bunches

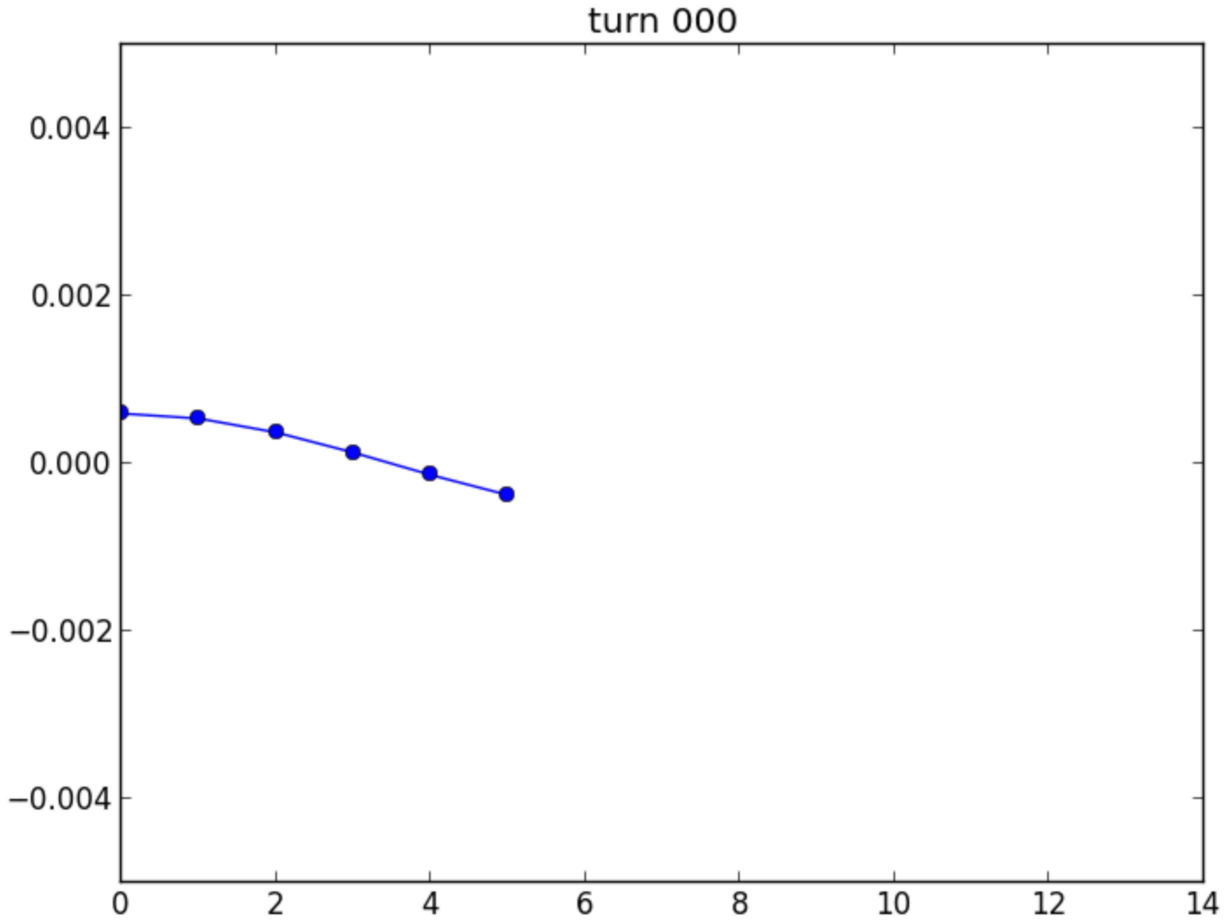

Why? Multi-bunch Instability

84 bunches

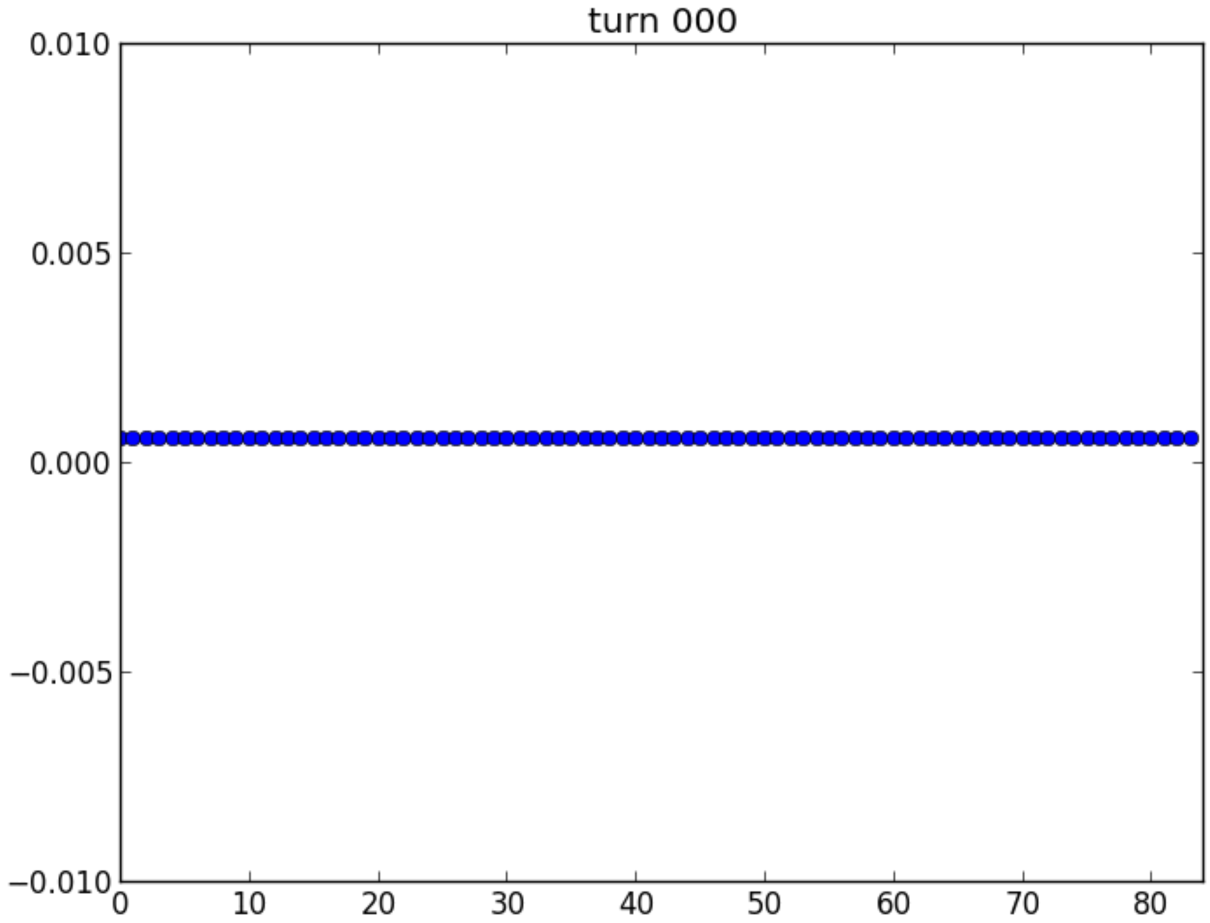

Why? Multi-bunch Instability

14 bunches

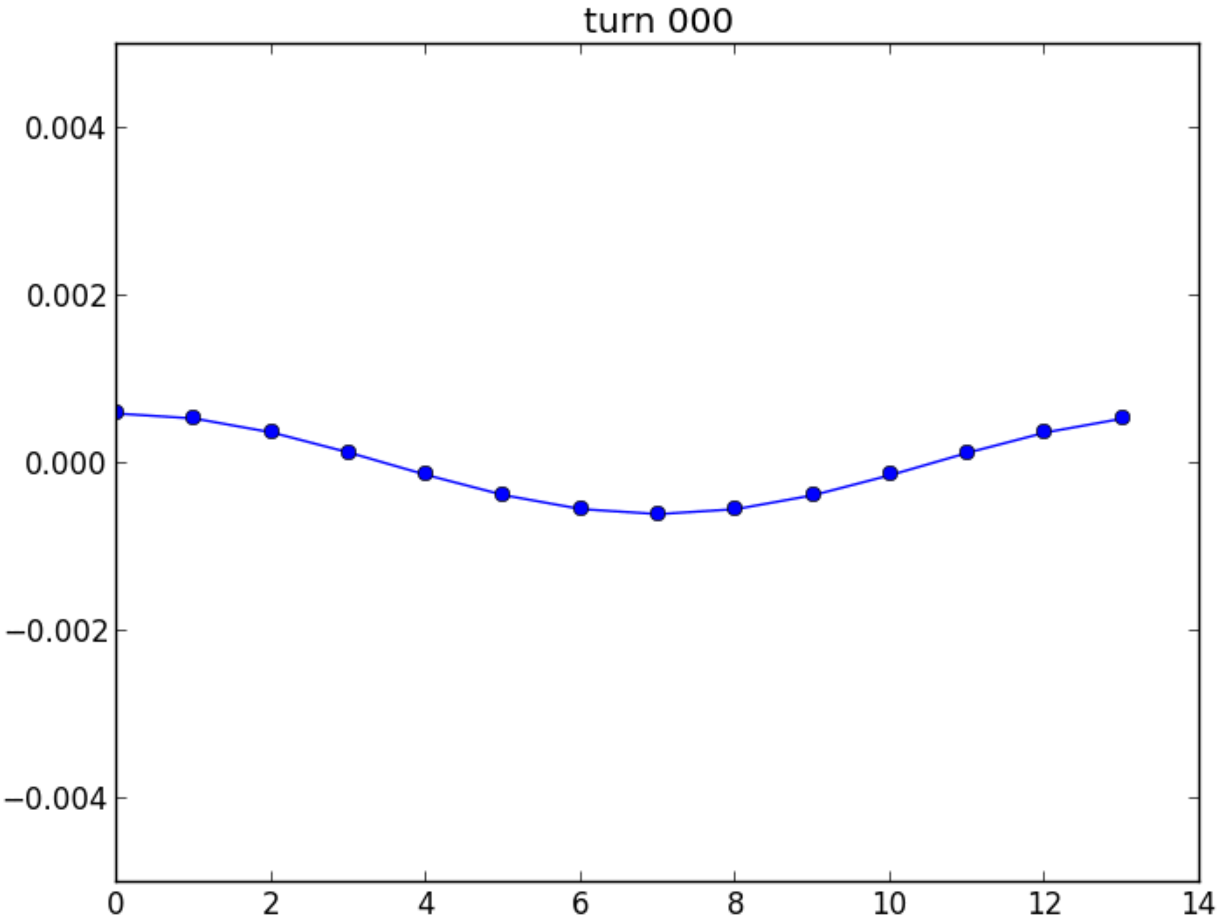

Application to Fermilab Intensity Frontier

- Slip stacking
  - Used at Fermilab to create high-intensity beams
    - Pairs of bunches combined
  - Synergia simulations of single pairs require O(1000) cores
    - Periodic boundary conditions mimic other pairs
  - Realistic simulations will include O(500) pairs
    - Non-trivial structure observed in operation
    - Bunch-bunch wake field interactions
- Truly a leadership class computing problem.
- Work in progress!

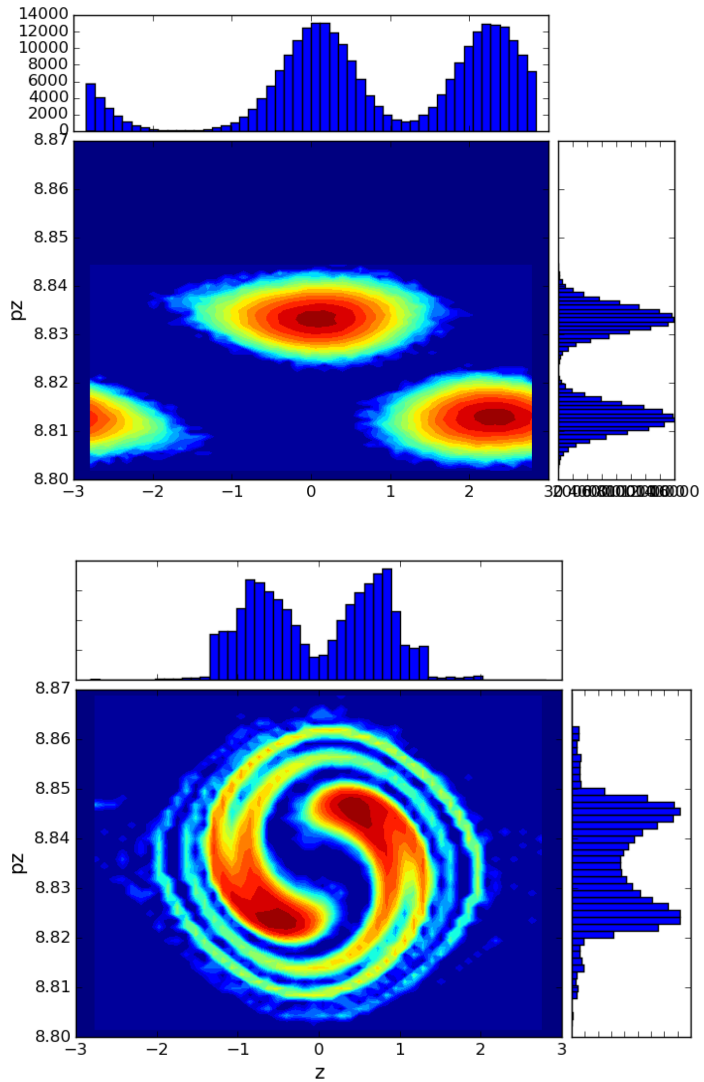

Fermilab

Application to Accelerator Theory

- Space charge modes provide theoretical framework for space charge studies
  - A. Burov, PRST-AB 12, 044202 (2009), PRST-AB 12, 109901, (2009).
- Difficult to extract modes from noise in realistic simulation
- First use of Dynamic Mode Decomposition (DMD) in Beam Dynamics
  - Macridin, et al., PRST-AB (2015).
- Excellent theory/simulation agreement

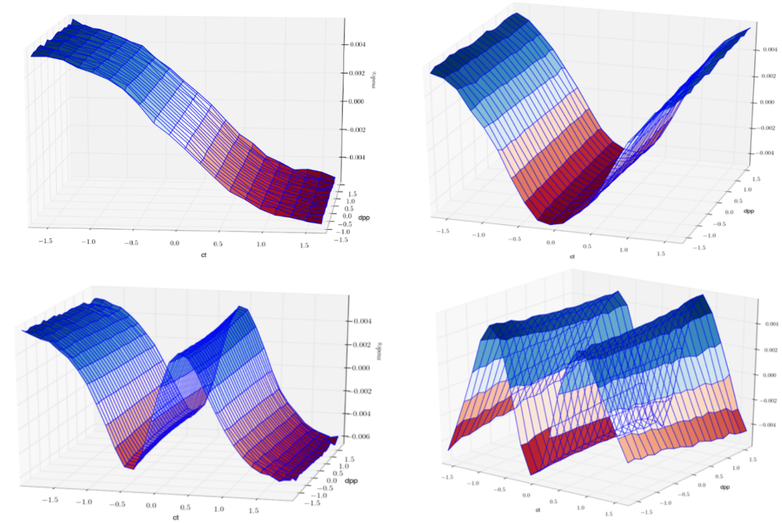

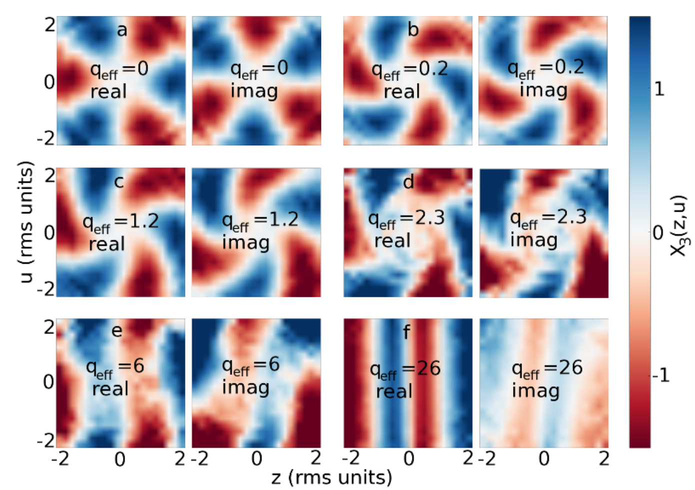

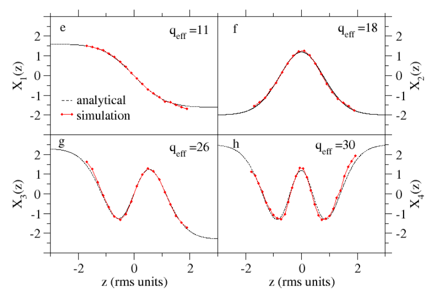

Application to High Luminosity LHC

- Study emittance growth over 100,000 revolutions in GSI SIS18 accelerator
  - Effects of statistical noise are important
- Largest beam dynamics simulation ever

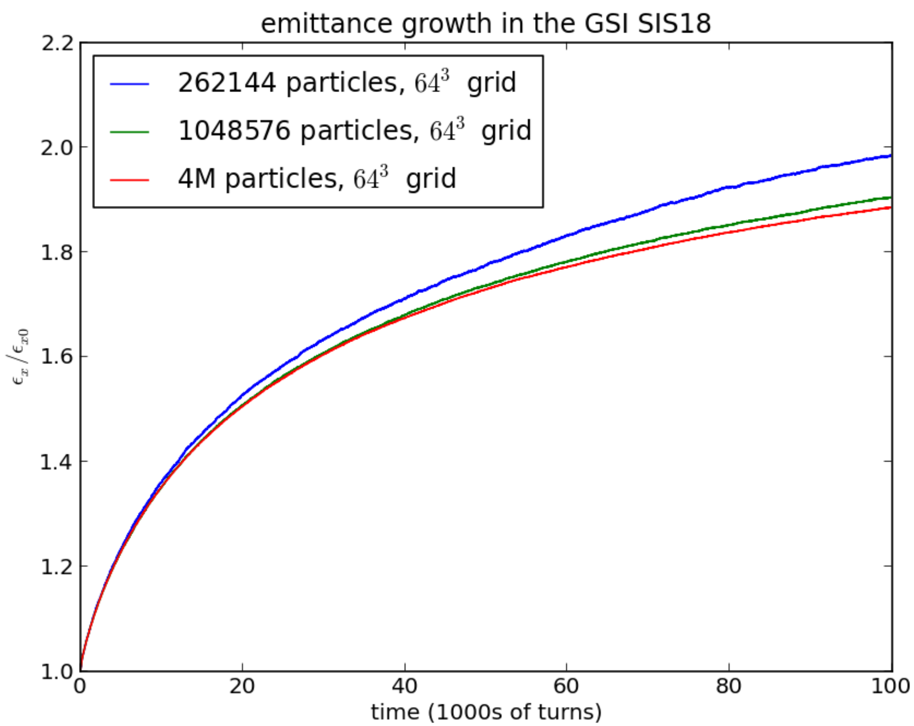

71 steps/turn
7,100,000 steps
4,194,304 particles
29,779,558,400,000 particle-steps
1,238,158,540,800,000 calls to "drift"

Yes, that's over a quadrillion

Fermilab

Next-generation Optimization

- Communication avoidance greatly optimized the collective portion of Synergia
- Single-particle performance is now critical
  - A quadrillion calls to drift!
  - High-statistics calculations increasingly important
- New technologies for HPC here and soon to be here
  - Multicore
  - GPU
  - Intel MIC
- General features
  - Memory/thread is decreasing
  - Vectorization is increasingly important (MIC has 8x SIMD)

Single-particle Calculations in Synergia

- CHEF (Collaborative Hierarchical Expansible Framework)
  - C++ library for single particles
    - Particles need to be converted to Synergia format
    - Abstraction penalty is small (few percent)
  - Complete implementation of single-particle physics
    - Uses same code to calculate particle trajectories and mappings
  - Dates to early 90's
- new: libFF (Fast Food)
  - Optimized for multi-particle calculation
    - particularly vectorization
  - Can use CHEF machinery to calculate mappings
  - More extensive use of higher-order techniques

Fermilab

Vectorization in Synergia

- Original data layout
  - Cache-friendly data locality
    - All coordinates for a single particle are contiguous
    - Not vectorization-friendly
  - Data stored in dense 2d array
    - Boost MultiArray
  - Independent particle code has per-particle overhead
    - Small*
    - Perfectly scalable

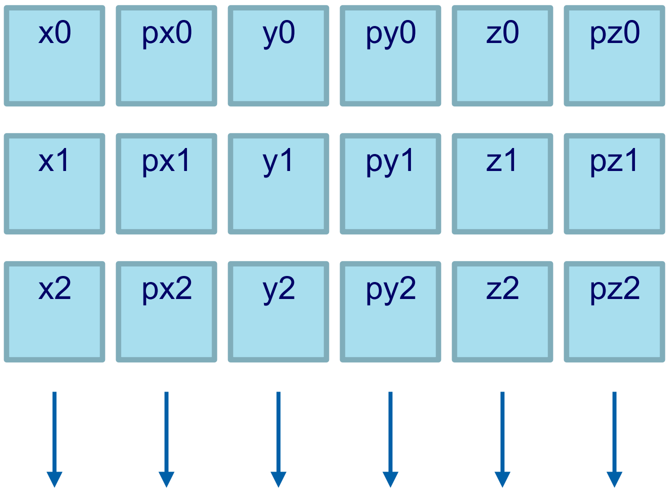

Fermilab

Vectorization in Synergia 2

- New data layout
  - Vectorization-friendly data locality
    - Each coordinate is contiguous
    - Always wins vs. original
  - Data still stored in dense 2d array
    - Boost MultiArray with Fortran ordering
      - Minimal code changes
  - New independent particle code has no per-particle overhead

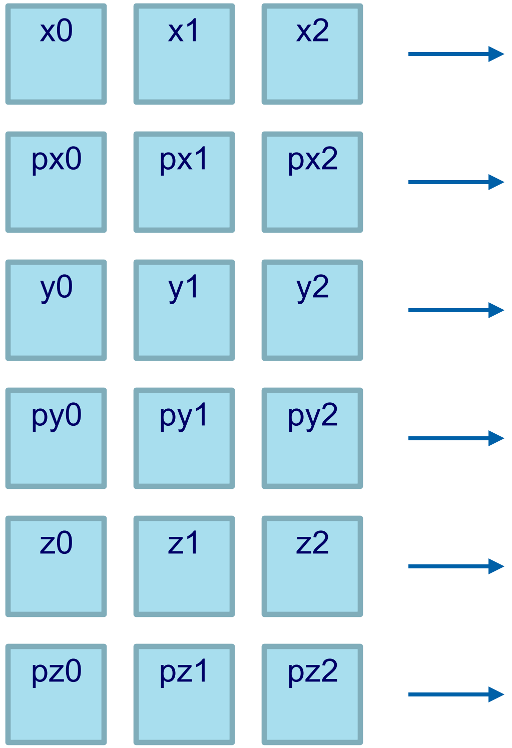

Explicit Vectorization

- C++ template-based
  - vectorclass
    - http://www.agner.org
  - GSVector
    - Generalized SIMD Vector
    - Part of Synergia
    - Compile-time vectorization model choice
      - double to AVX512
  - Already using templates for map calculations

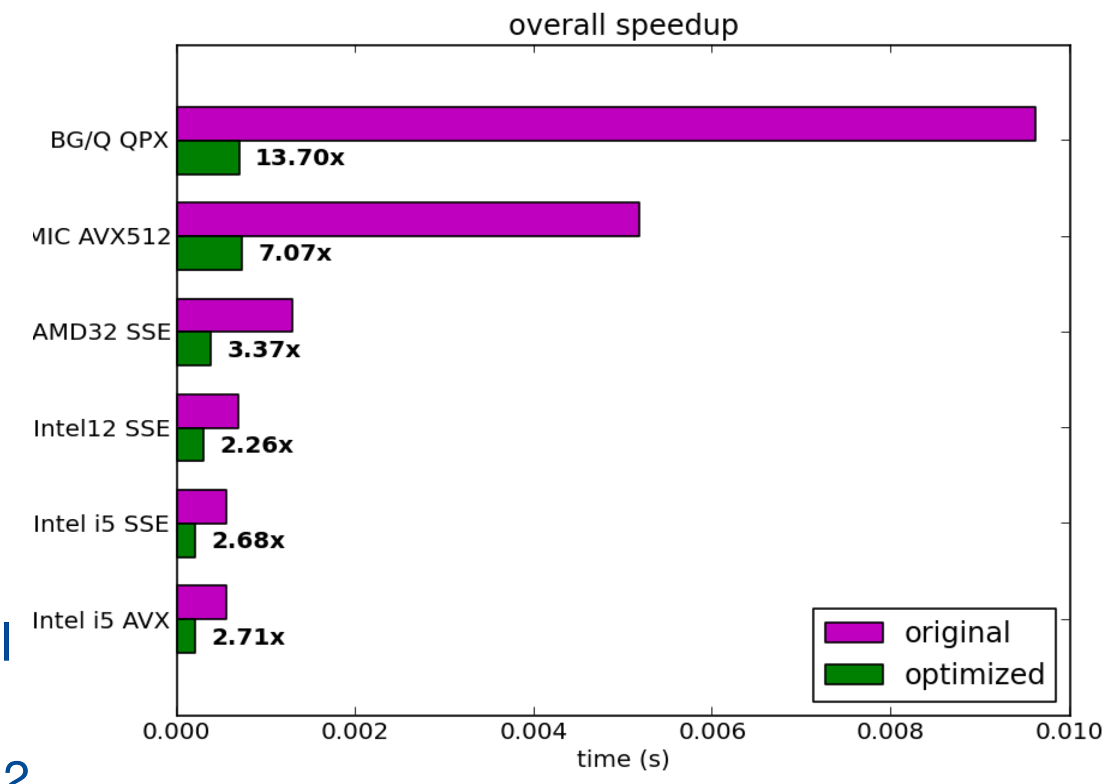

Fermilab

Multithread Optimization with OpenMP

- MPI does a good job of scaling single-particle physics
  – Will run out of memory
- OpenMP can be surprisingly difficult to scale
  – NUMA “first touch” needs to be taken into account
  – Reduction in the number of discrete processes greatly simplifies collective calculation

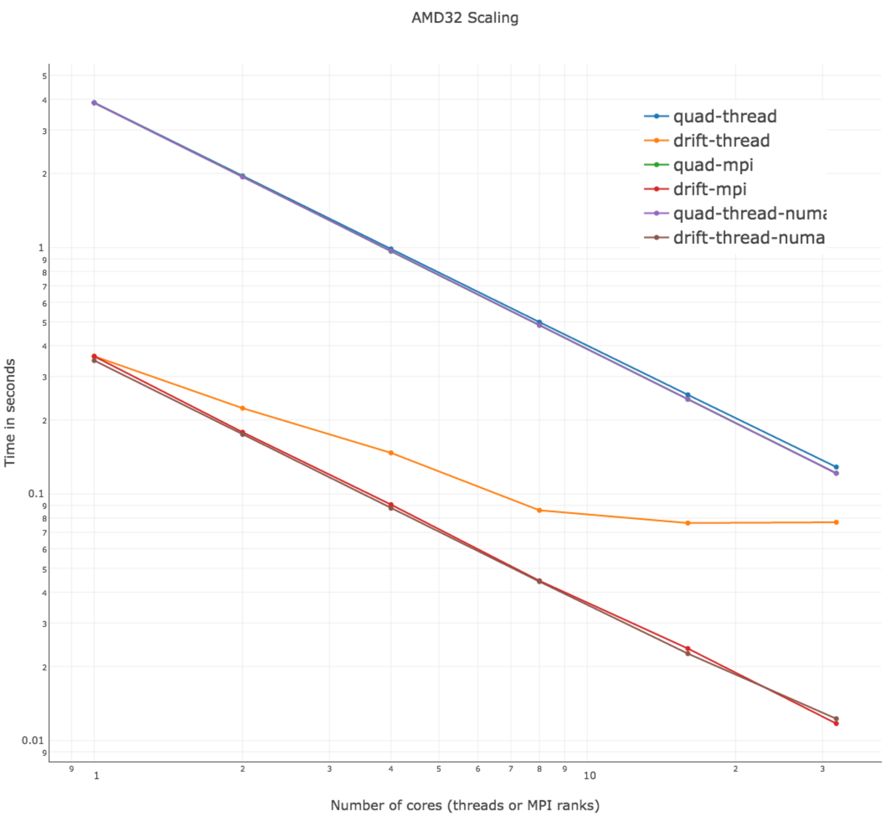

libFF + Explicit Vectorization + OpenMP on MIC

- libFF gives an overall speedup compared to CHEF
- Explicit vectorization gives another overall speedup
  – libFF required for vectorization
- OpenMP allows us to take advantage of all 240 hardware threads
- Overall speedup for quadrupole is 5768x

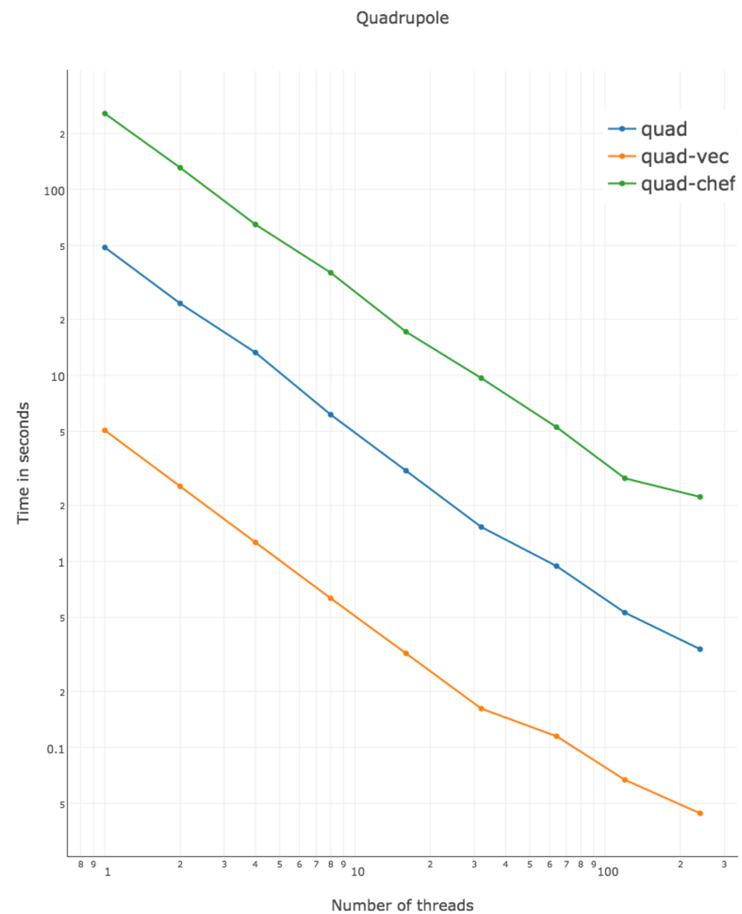

Improved Hierarchical Parallelism in Synergia Simulations

-1. Vectorization

0. Multithreading

1. Many particles/many grid points
2. Redundant solves within a bunch
   – Communication avoidance
3. Many bunches
4. Parameter scans/optimization

~~Four~~ Six levels of parallelism give us the potential to scale to a really huge number of parallel processes.

Conclusion

- Particle Physics depends on accelerators
  - Energy Frontier at CERN
  - Intensity Frontier at Fermilab
  - Higher intensities required for progress in both frontiers
- Synergia simulations are advancing our understanding of intensity-dependent effects in accelerators
  - Applications
    - Understanding instabilities in the Fermilab Booster
    - Modeling slip stacking in Fermilab Recycler and Main Injector
    - Evaluating theoretical models of space charge
    - Understanding long-term emittance growth for the High Luminosity LHC
      - Largest accelerator simulation ever
  - Evolving toward Exascale computing
    - Hierarchical parallelism with many levels